



Platform Notes: Sun™ Enterprise 6x00/5x00/4x00/3x00 Systems

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A. 650-960-1300

Part No. 806-4647-10
July 2000, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: (c) Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™: (c) Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Adobe PostScript

Contents

Preface v

1. OpenBoot Commands 1

OpenBoot Command Syntax 1

Environmental Monitoring Commands 2

Externally Initiated Reset XIR 3

▼ To Initiate an XIR 3

▼ To Display XIR Information 3

Hardware Watchdog 4

Flash PROM Management 4

System Configuration and Diagnostic Information 5

prtdiag(1M) Options 5

TOD Clock Management 5

NVRAM Variables 6

2. Board Hot-Plug Software and Dynamic Reconfiguration 9

Disabled System Board 9

▼ To Swap Out a Disabled Board 10

Activated System Board 10

▼ To Swap Out an Activated Type Three (PCI) Board 11

- ▼ To Swap Out an Activated Board if Your System Supports DR 11
 - Disabling Hardware 11
 - ▼ To Disable a Defective Board 12
 - Alternate Pathing 12
- 3. CPU Over-Temperature Safeguard 13**
 - COS Requirements 13
 - Factors in Overheating 14
 - COS Operation 14
 - Resolving an Over-Temperature Condition 14
 - ▼ To Resolve an Over-Temperature Condition 15
 - Failure to Disengage CPUs 15
 - Failure to Power Off CPUs 16

Preface

This book describes software features that apply only to the Sun Enterprise™ 6x00/5x00/4x00/3x00 family of servers.

Related Documents

For detailed information on the software features described in this book, refer to the man pages for Solaris 8. For information on DR features, refer to the *Sun Enterprise 6x00, 5x00, 4x00, and 3x00 Systems Dynamic Reconfiguration User's Guide*.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for SMCC Peripherals*
- AnswerBook™ online documentation for the Solaris™ software environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name</i> %
C shell superuser	<i>machine_name</i> #
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Alternate Pathing	<i>Sun Enterprise Server Alternate Pathing 2.3 Reference Manual</i>	806-1934
Dynamic reconfiguration	<i>Sun Enterprise 6x00, 5x00, 4x00, and 3x00 Systems Dynamic Reconfiguration User's Guide</i>	806-3984
OpenBoot 3.x	<i>OpenBoot 3.x Command Reference Manual</i>	806-1377

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

<http://www1.fatbrain.com/documentation/sun>

Accessing Sun Documentation Online

The `docs.sun.com`SM web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://docs.sun.com>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

docfeedback@sun.com

Please include the part number (806-4647-10) of your document in the subject line of your email.

OpenBoot Commands

This chapter describes the OpenBoot™ commands for Sun Enterprise 6x00/5x00/4x00/3x00 servers.

OpenBoot Command Syntax

This is an explanation of the syntax used in the examples of OpenBoot commands shown in the following sections.

In the examples a notation “(--)” represents the stack diagram, where any entries to the left of the “--” represent the state of the stack *before* the command is executed. Any any entries to the right represent the stack *after* the command is executed. If there are any such entries, the top of the stack is listed at the far right, preceded by entries that are lower in the stack. In the example below, the entries *c* and *z* are the top of the stack, before and after the command is executed, respectively.

```
command ( a b c -- x y z )
```

For more information about OpenBoot command usage, refer to OpenBoot documentation, such as the *OpenBoot 3.x Command Reference Manual*, part number 806-1377.

Environmental Monitoring Commands

The following commands are used for environmental monitoring.

TABLE 1-1 Environmental Monitoring Commands

Command	Purpose
<code>disable-environmental-monitor</code> <i>Usage:</i> <code>disable-environmental-monitor (--)</code>	To stop monitoring power supply status, board temperatures, and board hot plug while the screen displays the <code>ok</code> prompt.
<code>enable-environmental-monitor</code> <i>Usage:</i> <code>enable-environmental-monitor (--)</code>	To start monitoring power supply status, board temperatures, and board hot plug while the screen displays the <code>ok</code> prompt.

Externally Initiated Reset XIR

If a hard failure occurs, use XIR to reset the system, then immediately get information about the system state at the time of the hard failure.

▼ To Initiate an XIR

- **Use either the XIR button on the clock board or the remote console XIR sequence.**
When an XIR occurs, memory is cleared but some CPU states are temporarily saved. See the following section.

▼ To Display XIR Information

- **Type the following command at the ok prompt immediately after the XIR:**

```
ok .xir-state-all
```

The output displays the CPU state for each CPU:

TABLE 1-2 The XIR Display

```
#1 ok .xir-state-all
TL=1 TT=3
CPU ID#1
TPC=e0028688 TnPC=e0028688 TSTATE=9900001e06

CPU ID#5
TL=1 TT=3
TPC=e002755c TnPC=e0027560 TSTATE=4477001e03
#1 ok .xir-state-all
TL=1 TT=3
```

Where:

TL is the trap level.

TT is the trap type.

TPC is the trap program counter.

TState is the trap state.

Note – The `XIR` does not override the NVRAM `auto-boot?` variable.

Hardware Watchdog

The Sun Enterprise `xx00` family of servers provide the ability to enable a hardware timer that will hard-reset the system if it times out. To enable this feature, `watchdog_enable` must be set to 1 in the `/etc/system` file.

Flash PROM Management

The following OpenBoot 3.x commands are used in flash PROM management:

TABLE 1-3 Flash PROM Management Commands

Command	Purpose
<code>flash-update-system</code> <i>Usage:</i> <code>flash-update-system (--)</code>	To download the default flash images on all the boards in the system.
<code>prom-copy</code> <i>Usage:</i> <code>prom-copy (src dst --)</code>	To copy a flash PROM from board <code>src</code> to board <code>dst</code> . Source (<code>src</code>) and Destination (<code>dst</code>) are specified by slot number.
<code>update-proms</code> <i>Usage:</i> <code>update-proms (--)</code>	To synchronize the latest copy of each type of PROM on all other boards of the same type.

System Configuration and Diagnostic Information

To print system configuration and diagnostic information, use `prtdiag` instead of using an OpenBoot command. The diagnostic information displayed lists the failed field-replaceable units (FRUs) in the system. For more information on `prtdiag`, see the `prtdiag` man page.

`prtdiag(1M)` Options

- v lists the most recent AC Power failure, the most recent hardware fatal error information, and, if applicable, environmental status.
- l logs its output to `syslogd(1M)` only if failures or errors exist in the system.

TOD Clock Management

The following command is used for time-of-day (TOD) clock management, which includes NVRAM located on I/O boards.

TABLE 1-4 TOC Clock Management Command

Command	Purpose
<code>copy-clock-tod-to-io-boards</code> <i>Usage:</i> <code>copy-clock-tod-to-io-boards (--)</code>	To download the default flash images on all the boards in the system.

NVRAM Variables

The following table lists selected NVRAM variables and associated values that can be used with the dynamic reconfiguration (DR) feature on Sun Enterprise 6x00, 5x00, 4x00, and 3x00 systems:

TABLE 1-5 Specific NVRAM Variables

Variable Name	Comments
<code>configuration-policy</code>	This variable determines how the dynamic reconfiguration feature handles failing hardware. The value (<code>component</code> , <code>board</code> , or <code>system</code>) of the variable determines the level of hardware that is to be disabled if a failure occurs. The default value is <code>component</code> .
<code>disabled-board-list</code>	The value may be a list of hexadecimal slot numbers or it may be blank, meaning no boards are disabled. The default value is blank.
<code>memory-interleave</code>	This variable determines how memory is interleaved between CPU/Memory boards. The value may be <code>min</code> or <code>max</code> . Setting the value to <code>min</code> allows dynamic reconfiguration operations. Setting the value to <code>max</code> allows large memory configurations. The default value is <code>max</code> .
<code>sbus-probe-default</code>	The value may be a combination of numbers for the Sbus devices (0-3) and other on-board devices (hexidecimal <code>d</code> represents the on-board SOC) on an I/O board. Devices are probed in the order in which they are listed.
<code>sbus-specific-probe</code>	The value consists of a board number and a list of device numbers (0-3 and <code>d</code>), separated by a colon.

To set or unset these variables, use the `setenv` or `set-default` commands, as appropriate. Commands take effect on the next system reset or at power-on.

Here are some examples of command syntax:

```
ok setenv configuration-policy component
   (Disables only the component, if it fails.)
```

```
ok setenv configuration-policy board
   (Disables the entire board if a component fails.)
```

```
ok setenv disabled-board-list 7af  
(Disables boards in slots 7, 10 (hexidecimal a), and 15  
(hexidecimal f).)
```

```
ok set-default disabled-board-list  
(Removes all boards from the board list by setting it to null.)
```

```
ok setenv memory-interleave min  
(Turns off interleaving between boards.)
```

```
ok setenv sbus-probe-default 123d0  
(Sets order of probing to the second, third, and fourth SBus  
devices, followed by the SOC device and the first SBus device.)
```

```
ok setenv sbus-specific-probe 4:320  
(Specifies board 4 and SBus devices 3, 2, and 0, in order.)  
(Devices 1 and d are omitted and are therefore not probed.)
```


Board Hot-Plug Software and Dynamic Reconfiguration

This chapter explains how to use the board hot-plug capability and Dynamic Reconfiguration (DR), where applicable. Hot-plug is a hardware feature that allows the replacement of components while a server is running. However, the system cannot use the new component until the system is rebooted. The Dynamic Reconfiguration software enhancement allows replacement of certain I/O boards without needing to reboot the server.

For OpenBoot command information and examples, see Chapter 1.

The board hot-plug procedures described in this chapter need to be followed by a Dynamic Reconfiguration (DR) restart procedure or by a system reboot.

Note – For the Solaris 8 software environment, DR applies to Sun Enterprise type 1, 2, 4, and 5 I/O boards. These boards include the Sbus, Graphics, Sbus+, and Graphics+ I/O boards. I/O boards with slots for PCI cards are not yet supported.

Refer to the online *Sun Enterprise 6x00, 5x00, 4x00, 3x00 Systems Dynamic Reconfiguration User's Guide* for information about preparing your server to use DR, and instructions for performing DR operations. Check with your Sun Sales representative or system engineer for information on the availability of DR for your system.

Disabled System Board

A system board can be disabled (not used by the operating system) in three ways:

- A self-test detects a failure and disables the board.
- The board is disabled manually by using `disabled-board-list`. For example:

```
ok setenv disabled-board-list 72
```

The above command disables boards in slots 7 and 2.

- The board was inserted while the operating system was running.

▼ To Swap Out a Disabled Board

1. Verify that system precharge is OK.

Use the `prtdiag` command. For more information see “System Configuration and Diagnostic Information”.

2. Verify that hot-plug is available.

3. Make sure that the board is disabled (not in use by the operating system and the power light on the target board is off).

An example of the `setenv disabled-board-list` command is shown above.

4. Remove the disabled board.



Caution – If the yellow light is lit on the front panel, use `prtdiag` to determine the cause before installing a new board.

5. Install a new board.

6. To activate the board:

- For PCI boards, reboot the system.
- For other I/O boards, use DR commands.

Activated System Board

The system board is activated when the following three conditions are met:

- The board is in place during system startup.
- The slot has not been disabled by `disabled-board-list`.
- The board passes self-test.

▼ To Swap Out an Activated Type Three (PCI) Board

Dynamic reconfiguration does not yet support I/O boards with slots for PCI cards.

1. **Halt the system.**
2. **Power off the system.**
3. **Remove and replace the board.**
4. **Reboot the system to activate the board.**

▼ To Swap Out an Activated Board if Your System Supports DR

DR supports Sun Enterprise type 1, 2, 4, and 5 I/O boards. These boards include the Sbus, Graphics, Sbus+, and Graphics+ I/O boards.

1. **Use the appropriate `cfgadm` command to unconfigure the board from the system.**
Refer to the *Sun Enterprise 6x00, 5x00, 4x00, 3x00 Systems Dynamic Reconfiguration User's Guide*, part number 806-3984, for information about this command.
2. **Remove and replace the board.**
3. **To activate the board and reconfigure the board to the system, use DR commands.**

Disabling Hardware

By using the `configuration-policy` command, you can disable:

- A component—disables only the failing component
- A board—disables the board if any component on it fails
- A system—stops the system at the POST menu if there is a failure

Examples:

```
ok setenv configuration-policy board
```

or

```
# eeprom "configuration-policy=board"
```

▼ To Disable a Defective Board

If you suspect a board is defective and want to request that the system disable the board, perform the following procedure.

1. Isolate the board from the system:

- For a type three I/O board, use the `setenv` command at the `ok` prompt or the `eeprom` command at the `#` prompt and then reboot the system.
- For all other boards, use DR commands.

In the examples that follow, 3 means slot 3.

```
ok setenv disabled-board-list 3
```

or

```
# eeprom "disabled-board-list "=3
```

2. Remove the defective board and insert a new board.

3. To activate the board:

- For type 3 (PCI) I/O boards, clear the `disabled-board-list`, and then reboot the system.
- For all other boards, use DR commands.

See “Environmental Monitoring Commands” for more information.

Alternate Pathing

Alternate Pathing (AP) is a software enhancement that works in conjunction with DR and hot-plug. AP allows you to bypass a disabled disk or network adapter, avoiding unnecessary downtime. For additional information, refer to the Solaris on Sun Enterprise Servers AnswerBook module in the *Sun Enterprise Servers Alternate Pathing User's Guide*, part number 806-1934.

CPU Over-Temperature Safeguard

The CPU over-temperature safeguard (COS) is a Sun Enterprise xx00 platform feature for the Solaris 8 software environment and compatible versions available for servers with the proper firmware support. COS ensures that the temperature on any CPU/memory board *does not* exceed the safe operating range.

COS Requirements

COS is *not* available if a Sun Enterprise xx00 server lacks enabling firmware. In this case, the system displays the following messages during the boot sequence:

```
WARNING: Firmware does not support CPU power off
WARNING: Automatic CPU shutdown on over-temperature disabled
WARNING: Firmware does not support CPU restart from power off
WARNING: The ability to restart individual CPUs is disabled
```

When equipped with the proper firmware, the system displays the following during the boot sequence. Later firmware will show a similar output.

```
Board 0:  OBP  3.2.8 1997/02/27 14:00  POST 3.5.1 1997/03/05 09:34
```

- **To check the firmware revision level, use the `prtdiag -v` command.**

The correct firmware version for COS support is 3.2.8 or later.

Factors in Overheating

Many external conditions can raise the CPU/memory board temperature and compound high temperature problems, including:

- Room air-conditioning set incorrectly
- Lateral cooling obstructed

Some Solaris software environment issues can also affect the CPU temperature, such as bound threads or having only one CPU/memory board in the system. These Solaris software environment issues can cause a fallback to the existing shutdown behavior.

The CPU over-temperature safeguard does not affect the Solaris software environment in any way. COS operates only when the temperature of a CPU/memory board exceeds the safe operating range.

COS Operation

COS functions by monitoring the temperatures of all system CPUs. Warning messages are displayed in the system console if a CPU/memory board over-temperature condition occurs. The following example indicates an over-temperature condition for CPU/memory board 0:

```
WARNING: CPU/Memory board 0 is warm (temperature: 73C). Please
check system cooling
NOTICE: Processor 0 powered off.
NOTICE: Processor 1 powered off.
```

Resolving an Over-Temperature Condition

When the COS feature detects a CPU over-temperature condition, it takes the CPU offline and powers it off.

The system continues to operate with the offending CPU powered off. The CPUs are the chief source of heat on a CPU/Memory board; removing that heat source lowers the temperature into the normal operating range. This prevents sudden down time to the production server.

▼ To Resolve an Over-Temperature Condition

1. Verify the new state with the `psrinfo` command

The `psrinfo` output reflects the new CPU state:

```
0      powered-off since 03/11/97 09:48:31
1      powered-off since 03/11/97 09:48:31
```

2. Without powering off the operating system, replace the defective power supply (containing cooling fans) with a working unit.

Note – You can also halt the server using `/etc/halt` or `init 0` at the root or superuser prompt before replacing the defective power supply.

3. Bring the CPU back to normal operation using the `psradm` command:

```
# psradm -n processor_id#
```

With the CPU over-temperature safeguard feature, if the temperature sensor again reports an over-temperature (the temperature is still out of range), then the attempt to bring the CPUs back into operation using the `psradm` command fails, an exit status of `-1` and an error message is returned.

If the CPU in question has returned to normal operating temperature, the console displays a message similar to the following.

```
NOTICE: CPU/Memory board 0 has cooled down (temperature: 72C), system OK.
```

Failure to Disengage CPUs

In some instances, the CPU power control cannot disengage the affected CPU(s) from the Solaris software environment. For example, if the high temperature condition occurs when only one CPU/memory board with two processors is in the system, processor one will not go offline because it is the last processor in the system.

Failure to Power Off CPUs

If the attempted de-coupling of the problem CPU from the Solaris software environment fails, the temperature may continue to increase. When the temperature reaches the hard upper operational temperature limit, the system shuts down. In this case, a message similar to the following is displayed:

```
WARNING: CPU/Memory board 0 is very hot (temperature: 83C)
WARNING: System shutdown scheduled in 20 seconds due to over-temperature
condition on CPU/Memory board 0
WARNING: CPU/Memory board 0 still too hot (temperature: 83C). Overtemp shutdown
started
```